fragments can be used to create a goal to guide SECS toward such efficient syntheses, even though there may not be a reaction capable of doing that rejoining step.

Synthetic Analysis of Methyl Homodaphniphyllate: Being primarily concerned with the development of the SECS program, seldom does the opportunity arise to perform an extensive synthetic analysis on a particular molecule. At the beginning of our program to develop a sophisticated planning and strategy module we wanted to enumerate those things which chemists think about when planning a synthesis. By talking to other chemists and analyzing total syntheses which had been published in the literature we obtained a list of strategies which tells the chemist what to do. Of equal importance is a list of strategies which tells the chemist what not to do. In order to find these strategies we performed an extensive synthetic analysis on methyl homodaphniphyllate. The hope was that we would find useful generalizations that could be later used to prevent SECS from creating useless precursors. The compound used in the analysis was chosen because it is the sort of molecule which SECS handles best in its present form, that is, a molecule having few functional groups and a multi-bridged ring system. In addition, none of the alkaloids in this family, of which the present molecule is the simplest, have been synthesized.

The analysis of this material was carried through to depths of up to 14 levels and over six thousand precursors were generated. Several reasonable synthetic sequences emerged and some results of the analysis were reported at the Natural Products Symposium, part of the Western Regional ACS Meeting held in San Francisco on September 29, 1978. This analysis demonstrated the current capability of SECS with respect to very large problems. It further pointed out the great savings in time and effort that will result from even simple strategic control. This example serves as a base case to be compared with a later analysis employing more sophisticated strategic control.

Strategy Knowledge Base Building: Over the past year we have collected strategies, written them down, and searched for a uniform, formal method for representing these principles. To our knowledge, this is the first such thorough analysis of synthesis from this point of view, and it requires an effort similar to that for building a medical diagnosis knowledge base. Given such a knowledge base, our approach is to analyze the target molecule for problem areas. Each area may trigger certain pieces of knowledge which trigger others until finally goals are put on the goal list to direct SECS with respect to this particular problem area. We have studied many planning programs reported in the literature and have discovered that these programs strive to find one plan, for example, to cause a robot to accomplish a particular command. But we want not one plan, but all good plans for the synthesis. And as we expand the synthesis tree, the number of plans to be remembered increases. Thus the question arises of how to represent multiple plans. Our goal list essentially does that. By stating constraints that must be satisfied it excludes large regions of the tree. Thus one can think of this as a representation of all plans consistent with those constraints.

The example below shows a piece of knowledge relating to the control of stereochemistry.

```
IF 1) ATOM X IS A STEREOCENTER &
   2) ATOM X IS THE ORIGIN OF FG Y, 6
   3) STEREOGROUP Z IS WITHIN GAMA OF ATOM X ALONG PATH W &
   4) STERIC DIFFERENTIATION OF STEREOGROUP Z IS
       MEDIUM OR HIGH, &
   5) IT IS NECESSARY TO INCREASE THE STERIC
       DIFFERENTIATION OF ATOM X, THEN
CONCLUDE: STEREOSPECIFICALLY MIGRATE FG Y ALPHA TO Z
       ALONG PATH W. (0.95)
```

A principle based on symmetry states "It is useful to search for fragmentations such that one or more of the fragments have equivalent sites of attachment." Corey's synthesis of caryophyllene alcohol made use of this principle, although the pathway was not exactly that suggested directly by this principle. In our plans for next year we describe how we intend to use these principles.

Stereoisomer Generator: Our work with the SEMA stereochemical naming algorithm and application of the symmetry group of a chemical graph has led to a stereoisomer generator that has been tested on all possible cyclic saturated hydrocarbons having up to 15 atoms and 5 rings. The algorithm non-redundantly generates each stereoisomer, reports the symmetry group for that isomer, the canonical stereodescriptors, and then determines if the stereoisomer is chiral or achiral. Another module reports whether the structure is likely to be stable or not based on symbolic analysis of the ring system and stereochemistry. One potential application of this, besides simply enumerating stereoisomers, is to make it possible for a chemist to enter complex ring systems without specifying stereochemistry at obvious centers. This algorithm can then look at which of the possible stereoisomers are reasonable, and ask the chemist which he/she intended. This would relax the specification of stereochemistry to more nearly match normal chemical convention.

Metabolism Prediction: Numerous structurally different chemical compounds have been found to induce neoplasia in man and animals. In many cases these chemical carcinogens are metabolically activated by mammalian enzyme systems to their ultimate reactive and toxic structure. Many of the mechanisms involved in this "bioactivation" process are known or are in the process of being discovered. Thus, it is now possible based on the structure of a compound and a through knowledge of biotransformations to make rational predictions of the plausible metabolites of a compounds produced in a mammalian system. To study the metabolic activation of compounds we are creating a computer assistant which will generate the plausible metabolites of a compound utilizing the biotransformations known to occur in mammalian systems.

A new computer program called XENO for the metabolism of xenobiotic compounds has been developed based on technology from computer synthesis project. However, since metabolism is being simulated in the forward direction, whereas organic synthesis is simulated in the reverse direction, the XENO program is quite different in logic from SECS, although both use ALCHEM as a representation for reactions. The XENO data base of biotransforms was developed by careful survey of metabolism literature and consultation with a committee of metabolism

experts at NIH. We selected a mechanistic representation of metabolic processes
which means a small data base suffices to represent most of the known processes.
A critical evaluation of XENO by a panel of experts in Bethesda, Md. in February
1978 concluded that the data base of biotransforms must be considerably expanded,
but even now it is able to raise some interesting questions of alternative
metabolic pathways, etc. XENO is currently running on SUMEX-AIM.

   D. List of Current Project Publications

F. Choplin, R. Marc, G. Kaufmann, and W.T. Wipke, "Computer Design of Synthesis
     in Phosphorus Chemistry. Automatic Treatment of Stereochemistry," J. Chem.
     Info. and Computer Sci., 18, 110 (1978).

F. Choplin, R. Dorschner, G. Kaufmann, and W. T. Wipke, "Computer Graphics
     Determination and Display of Stereoisomers in Coordination Compounds," J.
     Organometallic Chem., 152, 101 (1978).

F. Choplin, C. Laurenco, R. Marc, G. Kaufmann, and W.T. Wipke, "Synthese Assistee
     par Ordinateur en Chimie des Composes Organophosphores," Nouveau J. de
     Chimie, 2, (3) 285 (1978).

W.T. Wipke, G. Ouchi, and S.Krishnan, "Simulation and Evaluation of Chemical
     Synthesis - SECS. An Application of Artificial Intelligence Techniques,"
     Artificial Intelligence, 10, 999 (1978).

M.L. Spann, K.C. Chu, W.T. Wipke, and G. Ouchi, "Use of Computerized Methods to
     Predict Metabolic Pathways and Metabolites," J. of Env. Pathology and
     Toxicology, 2, 123 (1978); also reprinted in "Hazards from Toxic
     Chemicals," ed. M.A. Mehlman, R.E. Shapiro, M.F. Cranmer and M.J. Norvell,
     Pathotox Publishers, Inc., Park Forest South, Ill., 1978, pp. 123-121.

In Press:

S.A. Godleski, P.v.R. Schleyer, E. Osawa, and W.T. Wipke, "The Systematic
     Prediction of the Most Stable Neutral Hydrocarbon Isomer," Progress in
     Physical Organic Chemistry, in press.

J.D. Andose, E.J.J. Grabowski, P. Gund, J.B. Rhodes, G.M. Smith, and W.T. Wipke,
     "Computer-Assisted Synthetic Analysis: The Merck Experience," in Computer-
     Assisted Drug Design, ACS Symposium Series, in press.

W.T. Wipke, D. Dolata, M. Huber, and C. Buse, "Machine Reasoning About
     Synthesis," in Computer-Assisted Drug Design, ACS Symposium Series, in
     press.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Collaborations and Medical Use of Programs via SUMEX.

SECS is available in the GUEST area of SUMEX for casual users, and in the SECS DEMO area for serious collaborators who plan to use a significant amount of time and need to save the synthesis tree generated. Much of the access by others has been through the terminal equipment at Santa Cruz because graphic terminals make it so much more convenient for structure input and output. We have assisted Professor J.E. McMurry of UCSC in his synthetic work towards aphidicoline and digitoxigenin (Total Synthesis of Cardiac Aglycones, HL-18118) using the model builder of SECS for evaluating plausible modes of ring closure. Numerous visitors to UC Santa Cruz have tried their own problems on the SECS program, generally taking away at least a couple of new ideas for research. Professor Ken Williamson of Mt. Holyoke College used SECS to build 3-D models of 50 compounds for C-13 nmr analysis, and his student provided us with a detailed report on their results and suggestions for improvements of our manual. Wilson Sallum of the University of Mass. Amherst working with Dr. E. McWhorter used SECS for the synthesis of various 3-naphthyl propionates. The synthesis suggested by SECS was successfully performed in the laboratory.

Synthetic chemists are beginning to come to us for a SECS analysis before beginning a laboratory synthesis. Dr. McMurry for example did a rather complete analysis of morphine before launching his recently successful synthesis. Plans for further new target analyses are underway between Dr. McMurry and Dr. Wipke.

Dr. Wipke has also used several SUMEX programs such as CONGEN in his course on Computers and Information Processing in Chemistry. Testing and collaboration on the XENO project with researchers at the NCI depend on having access through SUMEX and TYMNET.

B. Examples of Sharing, Contacts and Cross-fertilization with other SUMEX-AIM Projects:

We have had several discussions with the MYCIN group about our interest in an explanation capability for SECS. The AIM conference at Rutgers each year has been extremely valuable in generating ideas of new ways to apply current developments in AI to the problem of organic synthesis. Finally, it is impossible to count the daily exchanges that occur between researchers in the SECS group and other members of the AIM community on things related to languages, conferences, papers, seminars, and program sharing.

During the past year we have held weekly seminars on artificial intelligence related to the SECS project. These have been attended by Prof. Sharon Sickel (research area is theorem proving) and Prof. Michael Cunningham (research area: natural intelligence) of Information Sciences Dept. as well as

our group and other interested students and faculty. Visiting speakers include
Peter Friedland (Stanford MOLGEN project), Dennis Smith and Ray Carhart (both of
Stanford CONGEN project), Mark Stefik (Stanford MOLGEN), Jay Munyer (UCSC
analogical reasoning), Ken Friedenbach (UCSC and TRW, Hierarchical planning for
game of GO), and Stephan Unger (Syntex, drug design). This forum has been very
stimulating to our current research in strategies.

     John Kunz of the Pulmonary Function - Ventilator Management project
developed at UCSF utilizing SUMEX has requested and received a copy of INTERC.
This program was written to allow facile communication between the Santa Cruz
11/34 and SUMEX.

C. Critique of Resource Services:

     We find the SUMEX-AIM network very well human engineered and the staff very
friendly and helpful. The SECS project is probably one of the few on the AIM
network which must depend exclusively on remote computers, and we have been able
to work rather effectively via SUMEX. Basically we have found that SUMEX-AIM
provides a productive and scientifically stimulating environment and we are
thankful that we are able to access the resource and participate in its
activities. SUMEX-AIM gives us at UCSC, a small university, the advantages of a
larger group of colleagues, and interaction with people all over the country. We
especially thank SUMEX for support of the leased line for our GT40.

D. Collaborations and Medical Use of Programs via Computers other than SUMEX.

     Arrangements between the University of California, Santa Cruz and NIH have
been begun to try to install a version of SECS on the NIH PDP-10 computer system,
and possibly later on the NIH-CIS system. Under an arrangement approved in 1974
between First Data, Princeton University, and NIH, SECS has been available over
TELENET so that the public could evaluate the state of the technology first hand,
by simply contacting First Data. First Data was selected because that is the
system the NIH PROPHET program is also on. As a result of that arrangement,
anyone who wishes can use the SECS program without worrying about converting code
for their machine, and a number of people in the private sector both in the US
and abroad have done so. We are currently exploring updating the version of SECS
on ADP (First Data) and have recently installed a version on the University of
Penn. Medical School computer.

III. RESEARCH PLANS (7/79-7/81)

A. Long Range Project Goals and Plans.

     The SECS project now consists of two major efforts, computer synthesis and
metabolism, the latter being a very young project. Our plans for SECS for the
next year include adding a high level reasoning module for proposing strategies
and goals, and providing control which continues over several steps. This
reasoning module also will be able to trace the derivation of goals and thus
explain some of its reasoning. We also plan to focus on bringing the transform
library up in sophistication to improve the performance and capabilities of SECS.
Our library has been sufficient for previous testing, but now requires filling
gaps in its knowledge.

Currently the similarity module requires a special version of SECS. We plan in the next year to incorporate this module into the standard version of SECS so that the bonds that if broken could lead to identical or similar fragments can be used to create a goal to guide SECS toward such efficient syntheses, even though there may not be a reaction capable of doing that rejoining step. We still have not had an opportunity to improve the teletype interface which we hope to attack soon. Our hash coding scheme allows very rapid retrieval of compounds from libraries of compounds. It now remains to create appropriate data bases of available starting materials complete with stereochemistry and other technical data to enable us to explore some. starting-material oriented strategies. This will require an interactive data base builder/editor to be built first. Our users have brought to our attention ways to make SECS more machine independent, as well as suggestions for additions and improvements. We hope to assimilate these into our research goals wherever possible.

The XENO metabolism project will be expanding the data base to cover more metabolic transforms, including species differences, sequences of transforms, and stereochemical specificities of enzymatic systems. A second phase will apply our "similarity" function to determine when metabolites are similar to known carcinogens. We are also hoping to develop programs which will help maintain the growing data bases. It is not clear at this time how quantitative we can hope to be with XENO's predictions and that will be studied.

B. Justification and Requirements for Continued use of SUMEX.

The SECS and XENO projects require a large interactive time-sharing capability with high level languages and support programs. I am on the campus computing advisory committee and am the campus representative to the UC system-wide computing advisory committee and know that the UCSC campus is not likely in the future to be able to provide this kind of resource. Further there does not appear to be in the offing anywhere in the UC system a computer which would be able to offer the capabilities we need. Thus from a practical standpoint, the SECS and XENO projects still need access to SUMEX for survival. Scientifically, interaction with the SUMEX community is still extremely important to my research, and will continue to be so because of the direction and orientation of our projects. Collaborations on the metabolism project and the synthesis project need the networking capability of SUMEX-AIM, for we are and will continue to be interacting with synthetic chemists at distant sites and metabolism experts at the National Cancer Institute. Our requirements are for good support of FORTRAN. We now must run SECS overlayed, but the debugging tool DDT loses its symbol table during overlaying. This is a serious problem we hope can be fixed by SUMEX staff because without symbols, debugging is very difficult and time-consuming.

C. Needs beyond SUMEX-AIM.

Our needs are to develop local capabilities for printing, tape reading and writing. Our GT46 will be providing that this year. We also need some local production capability both to help offload SUMEX and to provide us needed computing when SUMEX is either not available or heavily loaded or load limited.

D. Recommendations for Community and Resource Development.

    The AIM workshop is excellent, particularly if it is held on the WEST COAST
once in a while. From a chemistry standpoint, the joint group meetings with the
DENDRAL group plus ability to attend seminars at Stanford and have visitors
participate in our seminar program really satisfy our needs for communication
with people of similar interests. We have proposed a workshop for the benefit of
the implementors rather than the principal investigators and administrators, for
that would do wonders to develop the human resource. We feel the computer
resource is rather efficiently used right now. The system does get sufficiently
busy that guests simply get almost no time and consequently decide the programs
they are using are poorly written and too slow. A system to handle guest
production would help both guests and researchers. Even for programs that are
still in research and development, some large scale testing is required which
resembles production and could benefit from this production machine. A trivial
suggestion but also important is that TV-EDIT be improved to not leave null
characters in files which cause problems with compilers both at SUMEX and at
other sites when the files are sent to another machine.

4.1.3   Hierarchical Models of Human Cognition

Hierarchical Models of Human Cognition (CLIPR Project)
Walter Kintsch and Peter G. Polson
University of Colorado
Boulder, Colorado

I. Summary of Research Program

The CLIPR project has only been on SUMEX since the first of the year, thus our work is in the earliest stages. However, one of our subgroups, the text comprehension group, has managed to accomplish a great deal during this time.

Technical Goals

The CLIPR project consists of two subprojects. The first, the text comprehension project,. is headed by Walter Kintsch and is a continuation of work on understanding of connected discourse that has been underway in Kintsch's laboratory for over seven years. The second, the planning project, is headed by Peter Polson of the University of Colorado and Michael Atwood of Science Applications Incorporated, Denver, and is studying the processes of planning using software design tasks.

The goal of the text comprehension project is to show how the components of the prose processing model described by Kintsch (1974; Kintsch and van Dijk, 1978) might be implemented in a HEARSAY-like control structure. Previous theoretical work has been oriented towards the development and evaluation of individual components of a global model of human prose processing. This work, as well as other research in cognitive psychology and artificial intelligence, has described a number of components necessary for a successful system. The current goal of this research is to describe the interaction of these components in the understanding of a segment of prose. We expect that the AGE formalism, of multiple independent, but cooperating knowledge sources way be a useful system in which to model such interactions. Thus, the primary task of the text understanding project is to make use of the theoretical tools that are provided on SUMEX to integrate and further extend Kintsch's theoretical and empirical work on the understanding of prose.

Similarly, the process of planning in complex domains like the design of software is conceptualized as involving the interaction of different kinds of knowledge at varying levels of abstraction. Skilled designers have extensive knowledge about the design process, as well as diverse knowledge of various algorithms and the constraints imposed by particular computer systems and programming languages. These pieces of knowledge must interact in complex ways in order to produce a detailed design for a given piece of software. We have assumed that a HEARSAY-like model of these interactions can adequately describe the design process and the planning mechanisms that underlie the construction of software designs. We plan to use AGE in order to model the planning process in the software design task and to construct simulations of protocols collected from experts doing actual designs.

E. A. Feigenbaum

Medical Relevance and Collaboration

The text comprehension project impacts indirectly on medicine, as the medical profession is no stranger to the problems of the information glut. By adding to the research on how computer systems might understand and summarize texts, and determining ways by which the readability of texts can be improved, medicine can only be helped by research on how people understand prose. Development of a more thorough understanding of the various processes responsible for different types of learning problems in children and the corresponding development of a successful remediation strategy would also be facilitated by an explicit theory of the normal comprehension process.

The planning project is attempting to gain understanding of the cognitive mechanisms involved in design and planning tasks. The knowledge gained in such research should be directly relevant to a better understanding of the processes involved in medical policy making and in the design of complex experiments. We are currently using the task of software design to describe the processes underlying more general planning mechanisms that are also used in a large number of task oriented environments like policy making.

Both the text comprehension project and the planning project involve the development of explicit models of complex cognitive processes; cognitive modelling is a stated goal of both SUMEX and research supported by NIMH.

The primary focus of collaborative activities for both CLIPR projects has involved interactions with Penny Nii and Edward Feigenbaum concerning the software tools needed to carry out our modelling activities. In addition, the text comprehension group has initiated some collaborative research with Alan Lesgold of the SUMEX SCP Project. This research involves the sharing of software tools developed by James Miller of the CLIPR project. Finally, SUMEX'S ARPANET facilities have enabled the sharing of information and research plans with Barbara and Frederick Hayes-Roth of the Rand Corporation, with whom the planning group's modelling efforts are being carried out.

Progress Summary

The bulk of the programming that has been done so far has been by the text comprehension group. A LISP program has been written to analyze a set of twenty texts and produce reasonable predictions of both the recall of information from and the readability of these texts. The first stage of this system is nearly completed, and preliminary reports of this work have already been presented (Kintsch, 1979).

The initial activities of the planning group have focused on the preliminary development of a theoretical model for the planning processes of software design and on learning about the software available at SUMEX for modelling this task.

Both groups have been involved in learning AGE, and how it can be applied to their individual domains.

List of Relevant Publications

Kintsch, W. and van Dijk, T. A. Toward a model of text comprehension and production. Psychological Review, 1978, 85, 363-394.

Atwood, M. E., Polson, P. G., Jeffries, R., and Ramsey, H. R. Planning as a process of synthesis. Technical Report SAI-78-144-DEN, Science Applications, Incorporated, Denver, Co. December, 1978.

Kintsch, W. On modelling comprehension. Invited address at the American Educational Research Association convention. San Francisco, April 10, 1979.

II.   Interactions with the SUMEX-AIM Resource

Sharing and Interactions with other SUMEX-AIM Projects

We have been working with Penny Nii and Edward Feigenbaum on the use of AGE as a modelling tool for both the prose comprehension project and the planning project. Feigenbaum and Nii have already made one 2-day visit to Colorado in which members of both projects were introduced to AGE. Access to theoretical tools like AGE are vital to the success of both projects.

The AGE super-structure will provide us a coherent framework within which
to articulate our ideas and will greatly reduce the resources required to develop
functioning models of comprehension and planning. In addition, by agreeing to
serve as trial users of this developing system we hope to provide useful input to
the AGE project staff. It is our hope that this collaboration will result in a
system that is truly useable for the development of complex models of cognitive
processes.

As noted above, the text comprehension project has discussed the
possibility of collaborative research with other SUMEX users. Alan Lesgold of
the Learning Research and Development Center at the University of Pittsburgh, a
member of the SUMEX SCP project, has expressed interest in the use of the prose
analysis program described above, as has James Voss of the Department of
Psychology at the University of Pittsburgh. We are considering the possibility
of making this program available to outside investigators via the guest facility
of "SUMEX.

Critique of Resource Management

The SUMEX-AIM resource is clearly suitable for the current and future needs
of our project. We have found the staff of SUMEX to be cooperative and effective
in dealing with special requirements and responding to our questions. The
facilities for communication on the ARPANET have also facilitated collaborative
work with investigators throughout the country.

III.   Research Plans (8/79 - 7/81)

Long Range Projects Goals and Plans

The long range plans of both CLIPR projects require extensive use of the
AGE facility as a basis for the development of the knowledge based systems that
we have described in preceding sections. The needs of the text understanding
project illustrate these requirements well. Although the prose program described
above generates reasonable predictions of recall and readability, certain aspects
of the predictions are clearly insufficient. These insufficiencies are caused by
the lack of real world knowledge in the procedures that are used to generate the
representation of the text. A more complete model of prose processing must be
able to access information ranging from word definitions to frame structures, and
we expect that AGE will be of use in the development of a model incorporating a
more adequate knowledge base. We also expect to make use of the UNITS package as
the basis for developing frame-like knowledge sources to be accessed by the AGE
control structure. Thus, the understanding project is dependent upon SUMEX
access in order to obtain both the necessary computing facilities and software
tools for the continued development of this work.

The primary goal of the planning project is the development of a model, or
a series of models, of human performance on the software design task. We intend
to begin by modeling the protocols of experts on a particular task, eventually
extending the model to other levels of experience and other tasks. To do this we
will have to become more familiar with AGE and work on articulating our theory in
a way that is compatible with the AGE framework. This will involve two parallel
lines of effort. One is a deeper analysis of our protocol data, to increase our

knowledge of the detailed planning processes and knowledge structures experts are using to solve these problems. The second is the development of a model in AGE that can simulate these processes. We have to date been using SUMEX only for the latter activity, but we are beginning discover that both objectives are so intertwined that it is counter-productive for us to be using separate computer systems. Thus we intend to transfer our protocol analysis activities to SUMEX. This will have the added advantage of making it easier for us to share this very rich data source with other investigators.

Justification and Requirements for Continued SUMEX Use

As noted in Section A, our research requires access to the AGE and UNITS systems, which are available only on SUMEX. In addition to any benefits we receive from access to SUMEX, AGE, and UNITS, the AGE and UNITS projects will also benefit from our testing of and experience with these experimental systems. Such interactions between the CLIPR and HPP projects have already been fruitful. We also expect that our interactions with Lesgold of the SCP project will continue by sharing both ideas and programs.

We anticipate that our CPU utilization may increase slightly due to the onset of our regular use of AGE. However, much of our programming efforts have been and will be isolated in non-peak early morning hours (due to the times in differences between Colorado and Stanford) and in overnight runs via the BATCH facility. Our CPU impact on everyday SUMEX use then will likely not increase.

In view of the additional files needed for AGE and UNITS, and the transfer of the planning group's protocols to SUMEX, our current disk allocation may become insufficient. We would thus appreciate an increase of 250 pages to a new total of 750 pages for our project. This increase, combined with use of the ARCHIVE facility for off-line storage of the majority of the planning group's protocols, should be sufficient for these needs.

Needs and Plans for Other Computational Resources

We currently use three other computing systems, two of which are local to the University of Colorado. One is the Department of Psychology's CLIPR system, which is a Xerox Sigma 3 used primarily for the real-time running of experiments to be modeled on SUMEX. The second is the University of Colorado's CDC 6400, which is used for various types of statistical analysis. Thirdly, the planning group has been using a PRIME computer located at Science Applications, Incorporated for the storage and analysis of protocols.

Being a remote site, we are clearly limited in our ability to get hard copy of SUMEX material, although the SUMEX staff has been most helpful in mailing whatever listings we need. We are now negotiating with the Boulder facility of the National Bureau of Standards for access to a PDP 11/40 that is connected to the ARPANET. This would provide us with hard copy in a way much more efficient for both ourselves and SUMEX. The tape drive on the 11/40 would also allow easier transfer of materials between SUMEX and our local computers.

Recommendations for Future Community and Resource Development

Our primary recommendation for future development within SUMEX involves (a)

the continued support of INTERLISP, which is needed for AGE and for other work we have underway on SUMEX and (b) the continued development of the AGE and UNITS projects. In particular we would like to see an extension of AGE to include a wider variety of control structures so that our psychological models would not be confined to one particular view of knowledge-based processing.

4.1.4   Higher Mental Functions Project

Higher Mental Functions Project

Kenneth Mark Colby, M.D.
Professor of Psychiatry and Computer Science
Neuropsychiatric Institute
University of California at Los Angeles

I. Summary of Research Program

A. Technical goals

The goals of this project are to contribute new knowledge and invention to the fields of psychiatry and neurology using concepts, methods and instruments of artificial intelligence. To achieve these goals, the project is involved in simulation studies of paranoid conditions, psychiatric taxonomy, and intelligent speech prostheses for patients with communication disorders.

B. Medical relevance and collaboration

The research has obvious medical relevance. The project collaborates with psychiatrists, neurologists, speech pathologists and neurolinguists.

C. Progress summary

During the past year the project has designed and constructed two intelligent speech prostheses, ISP-I and ISP-II. These devices consist of portable microprocessors and voice synthesizers. Part of the software consists of an orthographic-to-phonetic translator of several thousand rules and special cases written in the form of a production system. An ISP-I provides the user with an infinite vocabulary, error-corrective feedback, an ability to sound spell and the capacity for the user to create his own mnemonics for his own unique expressions. An ISP-I is designed for users who have not suffered central brain damage to the language system. Such users are patients with cerebral palsy, Parkinsonism, laryngectomy, and patients with tracheostomies in intensive care units.

An ISP-II, in addition to all the features of ISP-I, contains a lexical-semantic memory which is used to aid the word-finding problems of patients who have suffered brain damage. Such patients include strokes, brain tumors, and head traumas. The programs for these devices are first worked out and debugged on a big machine, the SUMEX facility, and then transferred to the microprocessors. Of particular help is the large English dictionary at SUMEX which we use both for the solution of orthographic-to-phonetic problems and for the organization of lexical memories to aid word-finding.

A few improvements have been made to the simulation of paranoia, PARRY, which now serves as an example to other research projects of how to go about simulating psychopathology.

The psychiatric classification scheme is unreliable in many respects. Hence this project has undertaken the task of trying to characterize patients according to their cognitive structures, properties in addition to conventional signs and symptoms. An algorithm which runs at SUMEX analyzes patient self-report accounts to find the conceptual patterns and key ideas underlying surface structure sentences. A profile of the patient is formed from the key ideas and patients with similar profiles are clustered into groups. This work is still in the exploratory pilot-study stage.

D. List of relevant publications

Colby, K. M. Mind Models: An Overview of Current Work. MATHEMATICAL
    BIOSCIENCES, 39, 159-185, 1978.

Colby, K. M., Christinaz, D. and Graham, S. A Computer-Driven, Personal,
    Portable, and Intelligent Speech Prosthesis. COMPUTERS AND BIOMEDICAL
    RESEARCH, 11, 337-343, 1978.

Colby, K. M., Faught, W. S., and Parkison, R. C. Cognitive Therapy of Paranoid
    Conditions: Heuristic Suggestions Based on a Computer Simulation Model.
    COGNITIVE THERAPY AND RESEARCH, 3, 55-60, 1979.

II. Interactions with the SUMEX-AIM Resource

A. Collaborations

As described above, this project uses SUMEX (1) to run PARRY (2) to write software for intelligent speech prostheses and (3) to construct a psychiatric taxonomy based on patients' cognitive structures.

B. Interactions with Other SUMEX-AIM Projects

The project interacts with other SUMEX projects at the University of Texas at Galveston and at Michigan State University.

C. Critique of resource management

Incredible as it may sound, we have no criticism of SUMEX, only praise. The members of our project uniformly agree SUMEX represents the best system we have ever worked with. The system is up almost all of the time, the personnel are cooperative and congenial, and suggested improvements are listened to and effected.

III. Research Plans

A. Long range project goals and plans

We plan to continue for the next two years to work on the above-described projects. If funding can be obtained, the taxonomy effort will be expanded into a full-scale effort.

B. Justification for SUMEX use

This project uses SUMEX for each of its research sub-projects as already described. We need a large machine that can run large LISP programs efficiently. We also need the large English dictionary available at SUMEX. No comparable facilities exist at UCLA. Hence we are quite dependent on SUMEX for the continuation of this research in psychiatry and neurology.

C. Other computational resources

Our other computational needs involve microprocessors and improved speech synthesizers. These can be constructed and developed in our laboratory at UCLA.

D. Recommendations

About once a month, an obscure bug appears in the ARPA net which shuts everything down. We would recommend this bug be discovered and dealt with mercilessly.

E. A. Feigenbaum

4.1.5   INTERNIST Project

INTERNIST Project

J. Myers, M.D. and H. Pople, Ph.D.
University of Pittsburgh
Pittsburgh, Pennsylvania

I.   Summary of Research Program

   A.   Technical Goals

      The major goal of the INTERNIST project is to produce a reliable and
adequately complete diagnostic consultative program in the field of internal
medicine. Although this consultative program is designed primarily to aid
skilled internists in complicated medical problems, the program may have spin-off
as a diagnostic and triage aid to physicians assistants, rural health clinics,
military medicine and space travel.

      To be effective, the program must be capable of multiple diagnoses (related
or independent) in a given patient and it should deal effectively with the time
axis in the development and course of disease states.

   B.   Medical Relevance and Collaboration

      The program inherently has direct and substantial medical relevance.

      The knowledge base should reach a critical stage of completeness within a
year, at which point we shall invite collaboration in the field testing of the
program in a number of medical institutions. Desires for such collaboration have
been very positively indicated by more than an adequate number of sister academic
health centers and community hospitals, etc.

      The Department of Pediatrics at Pittsburgh has engaged in a collaboration
with INTERNIST with the objective of a similar diagnostic program in the field of
pediatrics.

   C.   Progress Summary

      The original INTERNIST program described in previous progress reports and
documented in Pople, Myers & Miller [3] continues to be the standard diagnostic
program used to analyze clinical problems and to exercise newly developed
portions of the knowledge base.

      The structure of the medical knowledge base has remained comparatively
constant during the past year. The knowledge base has been expanded by the
addition of some sixty diseases plus twenty-nine in pediatrics. The existing
knowledge base is under a process of continual editing which attempts to keep the
data up to date by the addition new information about diseases as such becomes
available, and which expands and corrects the old data base as omissions or
errors are discovered. To our gratification, the progressive enlargement of the

knowledge base has in no significant adverse way affected the operation of the computer program.

The program and the knowledge base are continually being tested with challenging medical problems with good and reasonable success. The knowledge base remains too incomplete for any comprehensive or critical test on our hospital floors but the system is used on an ad hoc basis for clinical guidance.

Experience with this system has led to the identification of certain performance deficiencies that are being addressed in the design of a second generation diagnostic program (INTERNIST-II) the essential features of which are outlined in Pople [1]. A major objective in the design of the new program is to enable concurrent evaluation of the multiple components of a complex clinical problem, thereby enhancing the system's rate of convergence on the essential nature of the problem. A number of new concepts, not presently captured in the existing INTERNIST knowledge base, are required for this purpose; for example: the "constrictor" relation described in [1]; generalization of the INTERNIST disease hierarchy to a network permitting multiple categorization. For this purpose, a schema definition language has been devised, which enables the definition of disease categories such as "infectious disease," "collagen-vascular disease," "gastrointestinal hemorrhage," and others that cut across the basic INTERNIST hierarchy of organ system categories. Programs have been developed to map automatically onto these described nodes those terminal level disease entities which satisfy the node descriptions. By use of this expanded set of categories, the INTERNIST-II program is able to draw more precise boundaries around the sets of feasible hypotheses used to guide the acquisition of additional patient data. While still experimental, this new approach is expected to yield more efficient workup of complex clinical problems.

During 1978-79 two graduate students in computer science, one for the whole year and the other for six months, have made valuable contributions to INTERNIST both in the further development of the computer operating and analytical systems and in the organization and manipulation of the medical knowledge base.

One of our clinical fellows met an untimely death in November 1978 after contributing substantially to the medical knowledge base during his four months of activity. The other clinical fellow was diverted during the year from work in augmenting the medical knowledge base to the project of developing a CRT display and interface system for the clinical user of INTERNIST. This project became necessary because of the over 3,000 individual manifestations of disease in the system, which manifestations are necessarily arbitrarily worded at this point in development. The computer program utilized is ZOG, a very versatile menu selection system developed by Newell and colleagues at Carnegie-Mellon University. The project has been completed for our manifestations list as it exists today. It has proved to be very versatile and easy to use. A casual and new physician user can now learn in five minutes or so how to enter his data on a patient with a diagnostic problem and proceed to conclusions on the part of INTERNIST.

We underpredicted two matters involving time: (1) the many hours required to update and revise the existing medical knowledge base, and (2) the time required to program the necessary number of diseases required to bring the medical knowledge base to a "critical mass" for field testing. We are

approximately a year behind our original projected schedule. Nevertheless, real
progress has been made, to wit, the addition of some sixty new diseases and the
substantial revision of some previously programmed diseases. The continual
analysis of actual diagnostic problems in internal medicine has pointed out many
(in themselves) minor alterations needed in the knowledge base which, in the
composite, have provided for much smoother and more "intelligent" operation.

As of July 1, 1979, Doctor Randolph Miller, a previous junior collaborator
on INTERNIST, will have completed his formal graduate education in internal
medicine and will be joining the INTERNIST project as a full-time junior faculty
member (Assistant Professor of Medicine). Doctor Miller's presence and
contribution should allow, in collaboration with others working on the program,
the essential completing of the medical knowledge base in the academic year 1979-
80.

D. Publications

1. Pople, H.E. "The Formation of Composite Hypotheses in Diagnostic Problem
    Solving: An Exercise in Synthetic Reasoning", Proceedings of the Fifth
    International Joint Conference on Artificial Intelligence, Boston, August
    1977.

2. Pople, H.E. "On the Knowledge Acquisition Process in Applied A.I. Systems",
    Report of Panel on Applications of A.I., Proceedings of Fifth International
    Joint Conference on Artificial Intelligence, 1977.

3. Pople, H.E., Myers, J.D. & Miller, R.A. "The DIALOG Model of Diagnostic
    Logic and its Use in Internal Medicine, Proceedings of the Fourth
    International Joint Conference on Artificial Intelligence, Tbilisi, USSR,
    September 1975.

4. Pople, H.E. "Artificial Intelligence Approaches to Computer-based Medical
    Consultation, Proceeding IEEE Intercon, New York, 1975.

II. Interactions with SUMEX-AIM Resource

A, B. Collaborations and Medical Use of Program Via SUMEX

INTERNIST remains in a stage of research and development. As noted in the "Progress Summary" above, we are continuing to attempt to develop better computer programs to operate the diagnostic system, and the knowledge base cannot be used very effectively for collaborative purposes until it has reached a critical stage of completion. These factors have stifled collaboration via SUMEX up to this point and will continue to do so for the next year or two. In the meanwhile, through the SUMEX community there continues to be an exchange of information and states of progress. Such interactions particularly take place at the annual AIM Workshop.

Dr. Victor Yu, formerly associated with MYCIN, is now a faculty member at the University of Pittsburgh and has begun active participation in INTERNIST. Dr. Yu has been valuable in the programming of infectious diseases.

C. Critique of Resource Management

SUMEX has been an excellent resource for the development of INTERNIST. Our large program is handled efficiently, effectively and accurately. The staff at SUMEX have been uniformly supportive, cooperative, and innovative in connection with our project's needs.

III. Research Plans (8/78 to 7/81)

A. Long Range Project Goals and Plans

The primary goal of INTERNIST is to develop and complete an effective and reliable instrument for diagnostic consultation in internal medicine. To accomplish this a very extensive knowledge base must be developed, tested and continually updated. The initial stage of development is about 75% accomplished; a reasonable complete knowledge base, incorporating the new data structures identified in section I above, is a year in the future. With this development together with the improvement in the computer analytical program, INTERNIST will be suitable for a critical field trial, first in our own health center and, assuming success, in a half-dozen or so of additional health care institutions. Successful completion of the field test should make the program ready for practical clinical use.

B.   Justification and Requirements for SUMEX Use

     Neither the continued evaluation and development of INTERNIST's computer
program nor the manipulation and further development of INTERNIST's knowledge
base can be accomplished without a large computer resource such as SUMEX. SUMEX
has thus far met our requirements admirably and those requirements for the
research and development component of INTERNIST should remain relatively constant
over the next three years. The SUMEX resource (or its equivalent) is absolutely
essential to INTERNIST's progress.

C.   Needs and Plans for Other Computational Resources

     As predicted above, INTERNIST should be ready for field testing within two
years. It is realized that it is not the purpose to SUMEX in its present form to
support such extensive trials. Accordingly, a dedicated computer (or a dedicated
portion of SUMEX) will be needed to carry out the trials. No specific plans have
yet been made for this operation.

4.1.6   Medical Information Systems Laboratory

MISL - Medical Information Systems Laboratory

M. Goldberg, M.D. and B. McCormick, Ph.D.
University of Illinois at Chicago Circle

I. Summary of Research Program

Funding for the Medical Information Systems Laboratory (MISL) under NIH
grant 1-RO1-MB-00114 was terminated in the spring of 1978. While the Laboratory
continued its official existence for the last year, no active research was
conducted. Consequently, the Laboratory has not used SUMEX-AIM services.

II.   Interactions with the SUMEX-AIM resource

There has been no interaction to speak of between MISL and SUMEX-AIM.

III.   Research Plans

Part of the work begun under MISL has been continued under other projects.
Notably, continued development of a relational database system, RAIN, was funded
by the Defense Advanced Research Projects Agency. That work is now virtually
complete. It is expected that the Television Ophthalmoscopy (TVO) project,
funded by the National Eye Institute, will make use of the RAIN database system.
Now that RAIN is complete, TVO can proceed with its plans for an AI system called
STARE (for structured analysis of the retina). Continued access to SUMEX-AIM
would greatly benefit development of STARE, as it would facilitate communication
and possible collaboration with other researchers in the AI in medicine
community. It is hoped that the MISL account on SUMEX-AIM can be reassigned to
the TVO project.

E. A. Feigenbaum

4.1.7   PUFF/VM Project

        PUFF/VM: Biomedical Knowledge Engineering in Clinical Medicine

                            John J. Osborn, M.D.
                The Institutes of Medical Sciences (San Francisco)
                            Pacific Medical Center

                                    and

                        Edward A. Feigenbaum, Ph.D.
                        Computer Science Department
                            Stanford University

        The immediate goal of this project is the development of knowledge-based
programs to interpret physiological measurements made in clinical medicine. The
interpretations are intended to be used to aid in diagnostic decision making and
in therapeutic actions. The programs will operate within medical domains which
have well developed measurement technologies and reasonably well understood
procedures for interpretation of measured results. The programs are:

        (1) PUFF: the interpretation of standard pulmonary function laboratory data
             which include measured flows, lung volumes, pulmonary diffusion
             capacity and pulmonary mechanics, and

        (2) VM: management of respiratory insufficiency in the intensive care unit.

        The second, but equally important, goal of this project is the
dissemination of Artificial Intelligence techniques and methodologies to medical
communities that are involved in computer aided medical diagnosis and
interpretation of patient data.

I.   Summary Of Research Program

PUFF

    A. Technical Goals

        The task of PUFF program is to interpret standard measures of pulmonary
function. It is intended that PUFF produce a report for the patient record,
explaining the clinical significance of measured test results. PUFF also must
provide a diagnosis of the presence and severity of pulmonary disease in terms of
measured data, referral diagnosis, and patient characteristics. The program must
operate effectively over a wide range of pathological conditions with a broad
clinical perspective about the possible complexity of the pathology.

B. Medical Relevance and Collaboration

Interpretation of standard pulmonary function tests involves attempting to identify the presence of obstructive airways disease (OAD: indicated by reduced flow rates during forced exhalation), restrictive lung disease (RLD: indicated by reduced lung volumes), and alveolar-capillary diffusion defect (DD: indicated by reduced diffusivity of inhaled CO into the blood). Obstruction and restriction may exist concurrently, and the presence of one mediates the severity of the other. Obstruction of several types can exist. In the laboratory at the Pacific Medical Center (PMC), about 50 parameters are calculated from measurement of lung volumes, flow rates, and diffusion capacity. In addition to these measurements, the physician may also consider patient history and referral diagnosis in interpreting the test results and diagnosing the presence and severity of pulmonary disease.

Currently PUFF contains a set of about 60 physiologically based interpretation "rules". Each rule is of the form "IF <condition> THEN <conclusion>". Each rule relates physiological measurements or states to a conclusion about the physiological significance of the measurement or state.

The interpretation system operates in a batch mode, accepting input data and printing a report for each patient. The report includes: (1) Interpretation of the physiological meaning of the test results, the limitation on the interpretation because of bad or missing data; the response to bronchodilators if used; and the consistency of the findings and referral diagnosis. (2) clinical findings, including the applicability of the use of bronchodilators, the consistency of multiple indications for airway obstruction, the relation between test results, patient characteristics and referral diagnosis. (3) Interpretation Summary, which consists of the diagnosis of presence and severity of abnormality of pulmonary function.

C. Progress Summary

Knowledge base:

PUFF is implemented on the PDP-10 in a version of the MYCIN system which is designed to accept rules from new task domains. Currently approximately 60 pulmonary physiology rules related to the interpretation of measurements mentioned above have been implemented. A typical rule is:

If (FVC(PP)>=80) and (FEV1/FVC<predicted-5) then PEAK FLOW RATES ARE REDUCED, SUGGESTING AIRWAY OBSTRUCTION OF DEGREE
        if (predicted-15<= FEV1/FVC <predicted-5) MILD
        if (predicted-25<=FEV1/FVC <predicted-15) MODERATE
        if (predicted-35<=FEV1/FVC <predicted-25) MODERATE TO SEVERE
        if (FEV1/FVC <predicted-35)   SEVERE

Results

The results of the PUFF system are reviewed in more detail in the 1978 SUMEX annual report. A version of the PUFF system is now in routine daily use at

Pacific Medical Center. Reports are reviewed by a physician pulmonary physiologist. Over one half of the reports are accepted by the physician without change; they are signed and entered into the patient record. Most of the remaining reports are modified with the addition of a small point in the test interpretation.

During the past year, substantial progress has been made toward each of the goals identified in 1978, specifically PUFF was changed to:

(1) identify restrictive lung disease with greater accuracy.

(2) modify some of the existing rules on OAD,

(3) add rules to determine patient effort, or lack of effort, during the measurement acquisition,

(4) add rules related to blood gas analysis, and

(5) modify some parts of the PUFF program to increase the efficiency.

Table 1 summarizes agreement in severity of diagnoses made by two MD's and by PUFF rules. In 94% of 144 cases analyzed in a prospective study, the degree of severity (0=none; 1=mild; 2=moderate; 3=moderately-severe; 4=severe) of OAD diagnosed by the first MD was within a single degree of severity of OAD diagnosed by the second MD. In 96% of the 79 cases for which the first MD diagnosed OAD, the second MD diagnosed the severity of OAD within one level of the severity diagnosed by the first MD. Agreement within one degree of severity of the diagnoses by the first and second MD's was substantially lower in RLD and DD cases. These discrepancies occurred because the second MD consistently called RLD more severe than did the first MD, and he consistently did not diagnose diffusion defects when the first MD diagnosed DD of moderate or greater degree.

| | Percent Agreement with 1st MD | | | |
| | All 144 cases | | 1st MD made Dx | |
Diagnosis	Second M.D.	PUFF Rules	Second M.D.	PUFF Rules
Normal				
OAD	0.94	0.99	0.96	0.97
RLD	0.92	0.97	0.77	1.00
DD	0.87	0.87	0.60	0.80
Total	0.91	0.94	0.86	0.94

Table 1. Percent agreement within one degree of severity of diagnoses by two MD's and by the first MD and rules.